EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	0	(("US-5864837") or ("US-20020055966") or ("US-6389420") or ("US-5151989") or ("US-20030069902") or ("US-20020194206") or ("US-6321236") or ("US-5987506") or ("US-4897781") or ("US-6611895") or ("US-6173378") or ("US-6161169") or ("US-20020188821") or ("US-6182196") or ("US-6237067") or ("US-5657472") or ("US-5630096") or ("US-5123104") or ("US-20030093630") or ("US-6449700") or ("US-6154816") or ("US-5893160") or ("US-5860159")). PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/21 20:34
S2	48	(("5864837") or ("20020055966") or ("6389420") or ("5151989") or ("20030069902") or ("20020194206") or ("6321236") or ("5987506") or ("4897781") or ("6611895") or ("6173378") or ("6161169") or ("20020188821") or ("6182196") or ("6237067") or ("5657472") or ("5630096") or ("5123104") or ("20030093630") or ("6449700") or ("6154816") or ("5893160") or ("5860159")).PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/21 20:37
S3	23	S2 and queue	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 20:38
S4	12	S3 and parallel	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 20:39
S5	11	S4 and write	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 14:16

EAST Search History

S6	2	"20050108231"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 18:11
S7	2444	version near5 lock$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 18:11
S8	267142	read with write	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 18:12
S9	383	S7 and S8	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 18:14
S10	2435	"shared storage"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 18:14
S11	59	S9 and S10	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 18:14
S12	440	double adj1 linked	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:27
S13	8397	(707/8,200-206).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/22 20:28

EAST Search History

S14	21	S12 and S13	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:28
S15	15	S14 and @pd<"20031117"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:33
S16	4623	"binary tree"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:33
S17	21	S12 and S16	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:33
S19	13	S17 and @pd<"20031117"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:33
S20	2444	version near5 lock$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:45
S21	267142	read with write	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:45
S22	383	S20 and S21	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:45

EAST Search History

S23	2435	"shared storage"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:45
S24	59	S22 and S23	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:45
S25	39	S24 and @pd<"20031117"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/22 20:45

ACM PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
"pending transaction" +lock +write +read SEARCH
Feedback Report a problem Satisfaction survey

Terms used **pending transaction lock write read** Found **21** of **51 searched out of 51.**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 21 Result page: **1** 2

Relevance scale

1 Performance modeling of nested transactions in database systems
Hossam S. Hassanein, Mohamed E. El-Sharkawi
November 2000 **Proceedings of the 2000 conference of the Centre for Advanced Studies on Collaborative research**
Publisher: IBM Press
Full text available: pdf(148.66 KB) Additional Information: full citation, abstract, references, index terms

The nested transaction model was introduced to satisfy the requirements of advanced database applications. Moreover, it is currently the basic transaction model for new databases like workflow systems, mobile databases, and objectrelational databases. Though there are several performance evaluation studies of different concurrency control mechanisms in nested transactions, effects of transaction parameters on the overall system performance have not received any attention. In this paper, we study ...

Keywords: nested transactions, performance evaluation, simulation, two-phase locking

2 A new approach to developing and implementing eager database replication protocols
Bettina Kemme, Gustavo Alonso
September 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 3
Publisher: ACM Press
Full text available: pdf(449.43 KB) Additional Information: full citation, abstract, references, citings, index terms, review

Database replication is traditionally seen as a way to increase the availability and performance of distributed databases. Although a large number of protocols providing data consistency and fault-tolerance have been proposed, few of these ideas have ever been used in commercial products due to their complexity and performance implications. Instead, current products allow inconsistencies and often resort to centralized approaches which eliminates some of the advantages of replication. As an ...

Keywords: database replication, fault-tolerance, group communication, isolation levels, one-copy-serializability, replica control, total error multicast

3 Using histories to implement atomic objects
Tony P. Ng
November 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 4
Publisher: ACM Press
Full text available: pdf(2.74 MB) Additional Information: full citation, abstract, references, citings, index terms, review

In this paper we describe an approach to implementing atomicity. Atomicity requires that computations appear to be all-or-nothing and executed in a serialization order. The approach we describe has three characteristics. First, it utilizes the semantics of an application to improve concurrency. Second, it reduces the complexity of application-dependent synchronization code by analyzing the process of writing it. Third, our approach hides the protocol used to arrive at a ser ...

4 MIDDLE-R: Consistent database replication at the middleware level

Marta Patiño-Martinez, Ricardo Jiménez-Peris, Bettina Kemme, Gustavo Alonso
November 2005 **ACM Transactions on Computer Systems (TOCS)**, Volume 23 Issue 4
Publisher: ACM Press
Full text available: pdf(1.81 MB) Additional Information: full citation, abstract, references, index terms

The widespread use of clusters and Web farms has increased the importance of data replication. In this article, we show how to implement consistent and scalable data replication at the middleware level. We do this by combining transactional concurrency control with group communication primitives. The article presents different replication protocols, argues their correctness, describes their implementation as part of a generic middleware, Middle-R, and proves their feasibility with an extensive p ...

Keywords: Database replication, eager data replication, middleware, scalability

5 Cache coherence in systems with parallel communication channels & many processors

John C. Willis, Arthur C. Sanderson, Charles R. Hill
November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**
Publisher: IEEE Computer Society
Full text available: pdf(868.59 KB) Additional Information: full citation, abstract, references

This paper describes and analyzes two algorithms for maintaining cache coherence in multiprocessor systems with parallel communication channels and many processors. A distributed link-list relates all cache frames representing the same main memory block. Messages traverse the list to maintain list integrity, exclusive ownership, and consistent values. Memory access semantics are equivalent to a shared memory system without caches. Reference latency, efficiency of memory use, and hardware complex ...

6 The directory-based cache coherence protocol for the DASH multiprocessor

Daniel Lenoski, James Laudon, Kourosh Gharachorloo, Anoop Gupta, John Hennessy
May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture ISCA '90**, Volume 18 Issue 3a
Publisher: ACM Press
Full text available: pdf(1.74 MB) Additional Information: full citation, abstract, references, citings, index terms

DASH is a scalable shared-memory multiprocessor currently being developed at Stanford's Computer Systems Laboratory. The architecture consists of powerful processing nodes, each with a portion of the shared-memory, connected to a scalable interconnection network. A key feature of DASH is its distributed directory-based cache coherence protocol. Unlike traditional snoopy coherence protocols, the DASH protocol does not rely on broadcast; instead it uses point-to-point messages sent between th ...

7 Simplifying distributed database systems design by using a broadcast network

Jo-Mei Chang
June 1984 **ACM SIGMOD Record , Proceedings of the 1984 ACM SIGMOD international conference on Management of data SIGMOD '84**, Volume 14 Issue 2
Publisher: ACM Press
Full text available: pdf(1.36 MB) Additional Information: full citation, abstract, references, citings

Atomic broadcast and failure detection are powerful primitives for distributed database

systems In the distributed database system LAMBDA, they are provided as network primitives In this paper, we show how atomic broadcast and failure detection simplify transaction commitment, concurrency control, and crash recovery Specifically, we give a simple *two-phase non-blocking* commit protocol, whereas three phases are required in a point-to-point network We also give a simplified read-one/write-a ...

8 Dynamic voting algorithms for maintaining the consistency of a replicated database

S. Jajodia, David Mutchler
June 1990 **ACM Transactions on Database Systems (TODS)**, Volume 15 Issue 2
Publisher: ACM Press

Full text available: pdf(4.07 MB) Additional Information: full citation, abstract, references, citings, index terms, review

There are several replica control algorithms for managing replicated files in the face of network partitioning due to site or communication link failures. Pessimistic algorithms ensure consistency at the price of reduced availability; they permit at most one (distinguished) partition to process updates at any given time. The best known pessimistic algorithm, voting, is a "static" algorithm, meaning that all potential distinguished partitions can be listed in adv ...

9 Adaptive transaction scheduling

Abdelsalam Helal, Tung-Hui Ku, Ramez Elmasri, Sourav Mukherjee
December 1993 **Proceedings of the second international conference on Information and knowledge management**
Publisher: ACM Press
Full text available: pdf(985.56 KB) Additional Information: full citation, references, citings, index terms

10 Supporting cooperation in the Marvel process-centered SDE

Naser S. Barghouti
November 1992 **ACM SIGSOFT Software Engineering Notes , Proceedings of the fifth ACM SIGSOFT symposium on Software development environments SDE 5**, Volume 17 Issue 5
Publisher: ACM Press

Full text available: pdf(1.08 MB) Additional Information: full citation, abstract, references, citings, index terms

Large scale software development processes often require cooperation among multiple teams of developers. To support such processes, SDEs must allow developers to interleave their access to the various components of the projects. This interleaving can lead to interference, which may corrupt the project components. In traditional database systems, the problem is avoided by enforcing serializability among concurrent transactions. In traditional software development, the proble ...

11 Middleware for replication and transactions: Adaptive middleware for data replication

Jesús M. Milan-Franco, Ricardo Jiménez-Peris, Marta Patiño-Martínez, Bettina Kemme
October 2004 **Proceedings of the 5th ACM/IFIP/USENIX international conference on Middleware Middleware '04**
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(444.03 KB) Additional Information: full citation, abstract, references, citings

Dynamically adaptive systems sense their environment and adjust themselves to accommodate to changes in order to maximize performance. Depending on the type of change (e.g., modifications of the load, the type of workload, the available resources, the client distribution, etc.), different adjustments have to be made. Coordinating them is already difficult in a centralized system. Doing so in the currently prevalent component-based distributed systems is even more challenging. In this paper, we p ...

12 Shared memory synchronization: Transactional contention management as a non-

clairvoyant scheduling problem
Hagit Attiya, Leah Epstein, Hadas Shachnai, Tami Tamir
July 2006 **Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing PODC '06**
Publisher: ACM Press
Full text available: pdf(186.90 KB) Additional Information: full citation, abstract, references, index terms

The transactional approach to contention management guarantees atomicity by making sure that whenever two transactions have a conflict on a resource, only one of them proceeds. A major challenge in implementing this approach lies in guaranteeing progress, since transactions are often restarted.Inspired by the paradigm of *non-clairvoyant* job scheduling, we analyze the performance of a contention manager by comparison with an *optimal*, clairvoyant contention manager that knows the lis ...

Keywords: concurrency control, contention management, scheduling, software transactional memory, transactions

13 Isolation-only transactions for mobile computing
Qi Lu, M. Satyanaranyanan
April 1994 **ACM SIGOPS Operating Systems Review**, Volume 28 Issue 2
Publisher: ACM Press
Full text available: pdf(463.71 KB) Additional Information: full citation, citings, index terms

14 Sagas
Hector Garcia-Molina, Kenneth Salem
December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data SIGMOD '87**, Volume 16 Issue 3
Publisher: ACM Press
Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms

Long lived transactions (LLTs) hold on to database resources for relatively long periods of time, significantly delaying the termination of shorter and more common transactions. To alleviate these problems we propose the notion of a saga. A LLT is a saga if it can be written as a sequence of transactions that can be interleaved with other transactions. The database management system guarantees that either all the transactions in a saga are successfully completed or compensating transactions ...

15 The demarcation protocol: a technique for maintaining constraints in distributed database systems
Daniel Barbará-Millá, Hector Garcia-Molina
July 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 3
Publisher: Springer-Verlag New York, Inc.
Full text available: pdf(1.64 MB) Additional Information: full citation, abstract, references, citings

Traditional protocols for distributed database management have a high message overhead; restrain or lock access to resources during protocol execution; and may become impractical for some scenarios like real-time systems and very large distributed databases. In this article, we present the demarcation protocol; it overcomes these problems by using explicit consistency constraints as the correctness criteria. The method establishes safe limits as "lines drawn in the sand" for updates, and makes i ...

Keywords: consistency constraints, serializability, transaction limits

16 On-chip communication architectures: analysis and optimisation: Power analysis of

system-level on-chip communication architectures
Kanishka Lahiri, Anand Raghunathan
September 2004 **Proceedings of the 2nd IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis**
Publisher: ACM Press
Full text available: pdf(101.16 KB) Additional Information: full citation, abstract, references, index terms

For complex System-on-chips (SoCs) fabricated in nanometer technologies, the system-level on-chip communication architecture is emerging as a significant source of power consumption. Managing and optimizing this important component of SoC power requires a detailed understanding of the characteristics of its power consumption.Various power estimation and low-power design techniques have been proposed for the global interconnects that form part of SoC communication architectures (*e.g.*, low- ...

Keywords: communication architectures, low-power design, network-on-chip, power analysis, system-on-chip

17 P2P systems & data streams: Reconciling while tolerating disagreement in collaborative data sharing
Nicholas E. Taylor, Zachary G. Ives
June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**
Publisher: ACM Press
Full text available: pdf(400.54 KB) Additional Information: full citation, abstract, references, index terms

In many data sharing settings, such as within the biological and biomedical communities, global data consistency is not always attainable: different sites' data may be dirty, uncertain, or even controversial. Collaborators are willing to *share* their data, and in many cases they also want to selectively *import* data from others --- but must occasionally *diverge* when they disagree about uncertain or controversial facts or values. For this reason, traditional data sharing and da ...

Keywords: collaborative data sharing, data integration, peer-to-peer, reconciliation, transactions, updates

18 Session 1: A recovery algorithm for a distributed database system
Nathan Goodman, Dale Skeen, Arvola Chan, Umeshwar Dayal, Stephen Fox, Daniel Ries
March 1983 **Proceedings of the 2nd ACM SIGACT-SIGMOD symposium on Principles of database systems**
Publisher: ACM Press
Full text available: pdf(823.78 KB) Additional Information: full citation, abstract, references, citings

We describe a reliability algorithm being considered for DDM, a distributed database system under development at Computer Corporation of America. The algorithm is designed to tolerate clean site failures in which sites simply stop running. The algorithm allows the system to reconfigure itself to run correctly as sites fail and recover. The algorithm solves the subproblems of atomic commit and replicated data handling in an integrated manner.

19 Are crossbars really dead?: the case for optical multiprocessor interconnect systems
Andreas G. Nowatzyk, Paul R. Prucnal
May 1995 **ACM SIGARCH Computer Architecture News , Proceedings of the 22nd annual international symposium on Computer architecture ISCA '95**, Volume 23 Issue 2
Publisher: ACM Press
Full text available: pdf(1.16 MB) Additional Information: full citation, abstract, references, citings, index terms

Crossbar switches are rarely considered for large, scalable multiprocessor interconnect systems because they require $O(n^2)$ switching elements, are difficult to control efficiently

and are hard to implement once their size becomes too large to fit on one integrated circuit. However these problems are technology dependent and a recent innovation in fiber optic devices has led to a new implementation of crossbar switches that does not share these problems while retaining the full advanta ...

20 Logical design of a reliable transaction management in a distributed multiple processor system

Herbert Kuss

December 1983 **Proceedings of the 1983 ACM SIGSMALL symposium on Personal and small computers**

Publisher: ACM Press

Full text available: pdf(648.14 KB) Additional Information: full citation, abstract, references, index terms

The logical design of a reliable transaction management in a distributed multiple processor system is presented. A main characteristic of the system is that it realizes autonomous nodes. This has implications on the construction of the distributed data base management software. We discuss some construction problems and we provide the concept of guardian transaction managers for the implementation of transactions as atomic units of consistency and recovery.

Results 1 - 20 of 21 Result page: **1** 2

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

ACM PORTAL

USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

"pending transaction" +lock +write +read

SEARCH

Feedback Report a problem Satisfaction survey

Terms used **pending transaction lock write read**

Found **21** of **51**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 21 - 21 of 21 Result page: previous 1 **2**

Relevance scale

21 Experiences in functional validation of a high level synthesis system

Ranga Vemuri, Paddy Mamtora, Praveen Sinha, Nand Kumar, Jay Roy, Raghu Vutukuru

July 1993 **Proceedings of the 30th international conference on Design automation**

Publisher: ACM Press

Full text available: pdf(911.68 KB) Additional Information: full citation, references, citings, index terms

Results 21 - 21 of 21 Result page: previous 1 **2**

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

IEEE Xplore RELEASE 2.1

Home | Login | Logout | Access Information | Alerts | Sitemap | Help

Welcome United States Patent and Trademark Office

Search Results BROWSE SEARCH IEEE XPLORE GUIDE SUPPORT

Results for "(('execution queue')<in>metadata)"

e-mail printer friendly

Your search matched **1** of **1415139** documents.

A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

» Search Options

View Session History

New Search

Modify Search

(('execution queue')<in>metadata) Search

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

view selected items Select All Deselect All

1. **A decentralized task scheduling algorithm and its performance modeling for computer networks**
 Ahmad, I.; Ghafoor, A.; Mehrotra, K.;
 Parallel and Distributed Processing, 1991. Proceedings of the Third IEEE Symposium on
 2-5 Dec. 1991 Page(s):314 - 321
 Digital Object Identifier 10.1109/SPDP.1991.218225
 AbstractPlus | Full Text: PDF(676 KB) IEEE CNF
 Rights and Permissions

Help Contact Us Privacy & Security IEEE.org

Indexed by Inspec

© Copyright 2006 IEEE – All Rights Reserved